**UNITED STATES DISTRICT COURT**
**IN THE EASTERN DISTRICT OF TEXAS**
**TYLER DIVISION**

| | |
|---|---|
| DATA SCAPE LIMITED,<br><br>Plaintiff,<br><br>v.<br><br>DELL TECHNOLOGIES INC., DELL INC., and EMC CORPORATION,<br><br>Defendants. | C.A. No. 6:18-cv-00658<br><br>**JURY TRIAL DEMANDED** |

## COMPLAINT FOR PATENT INFRINGEMENT

This is an action for patent infringement arising under the Patent Laws of the United States of America, 35 U.S.C. § 1 *et seq.* in which plaintiff Data Scape Limited ("Plaintiff," "Data Scape") makes the following allegations against defendants Dell Technologies Inc., Dell Inc., and EMC Corporation (each "Defendant" or collectively "Defendants"):

## PARTIES

1. Data Scape is a company organized under the laws of Ireland with its office located at Office 115, 4-5 Burton Hall Road, Sandyford, Dublin 18, Ireland.

2. On information and belief, defendant Dell Technologies Inc. is a Delaware corporation with a principal place of business at One Dell Way, Round Rock, Texas 78682. Data Scape is informed and believes that Dell Technologies has additional offices at 2300 West Plano Parkway, Plano, Texas. Dell Technologies may be served through its registered agent, Corporation Service Company, 251 Little Falls Drive, Wilmington, DE 19808.

3. On information and belief, defendant Dell Inc. is a Delaware corporation with a principal place of business at One Dell Way, Round Rock, Texas 78682. Dell. Inc. has additional offices at 2300 West Plano Parkway, Plano, Texas. Dell Inc. is wholly owned by its corporate parent, Dell Technologies Inc. Dell may be served through its

registered agent, Corporation Service Company, 251 Little Falls Drive, Wilmington, DE 19808.

4. On information and belief, defendant EMC Corporation is a Massachusetts corporation with a principal place of business at One Dell Way, Round Rock, Texas 78682, and a secondary place of business at 176 South Street, Hopkinton, Massachusetts 01748. EMC Corporation has further business operations, including a test-laboratory, at 3801 E. Plano Parkway, Suite 150, Plano, Texas. EMC may be served through its registered agent, Corporation Service Company, at 211 E. 7th St., Austin, TX. EMC Corporation is wholly owned by its corporate parent, Dell Technologies Inc.

## JURISDICTION AND VENUE

5. This action arises under the patent laws of the United States, Title 35 of the United States Code. This Court has original subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

6. This Court has personal jurisdiction over each defendant in this action because each defendant has committed acts within the Eastern District of Texas giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over each defendant would not offend traditional notions of fair play and substantial justice. Each defendant, directly and through subsidiaries or intermediaries, has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the asserted patents.

7. Venue is proper in this district under 28 U.S.C. § 1400(b). Each defendant has established places of business in the Eastern District of Texas. Each defendant is registered to do business in Texas. Upon information and belief, each defendant has transacted business in this District and has committed acts of direct and indirect infringement in this District.

## COUNT I

## INFRINGEMENT OF U.S. PATENT NO. 8,386,581

8. Data Scape is the owner by assignment of United States Patent No. 8,386,581 ("the '581 Patent"), entitled "Communication System And Its Method and Communication Apparatus And Its Method." The '581 Patent was duly and legally issued by the United States Patent and Trademark Office on February 26, 2013. A true and correct copy of the '581 Patent is included as Exhibit A.

9. Each defendant has offered for sale, sold and/or imported into the United States products and services that infringe the '581 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, Defendant's products and services, *e.g.*, Data Domain software and hardware, RecoverPoint software and hardware, and all versions and variations thereof since the issuance of the '581 Patent ("Accused Instrumentalities").

10. Each defendant has directly infringed and continues to infringe the '581 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities. Each defendant uses the Accused Instrumentalities for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

11. For example, the Accused Instrumentalities infringe Claim 1 (as well as other claims) of the '581 Patent. One non-limiting example of the Accused Instrumentalities' infringement is presented below:

12. The Accused Instrumentalities include "[a] communication apparatus." For example, the Accused Instrumentalities communicate data stored on one device (e.g. a Data Domain system) to another device (e.g. a second Data Domain system). *See, e.g.*, EMC Data Domain Operating System Version 5.7 Administration Guide ("Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of

management. *** The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. A Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination for one or more pairs. After replication is started, the source system automatically sends any new backup data to the destination system.").

13. The Accused Instrumentalities include "a storage unit configured to store content data to a storage medium" For example, a Data Domain system includes disks and/or solid-state storage medium. *See, e.g.*, Administration Guide ("Storage in most Data Domain systems is set up in a double parity RAID 6 configuration (two parity drives). Additionally, most configurations include a hot spare in each enclosure, except the DD1xx series systems, which use eight disks. *** To keep data synchronized during a hardware or power failure, the Data Domain system uses NVRAM (non-volatile RAM) to track outstanding I/O operations. An NVRAM card with fully charged batteries (the typical state) can retain data for a period of hours, which is determined by the hardware in use."). The second storage medium is configured to store management information of data to be transferred, e.g. replication configuration settings, folder metadata, etc.

14. The Accused Instrumentalities further include "a communication unit configured to communicate with an external apparatus." For example, a Data Domain system can connect to another Data Domain system over a wide area network. *See, e.g.*, Whitepaper H7082: Dell EMC Data Domain Replicator ("In comparison, replication uses the wide area network (WAN) as the transport mechanism for data instead of tapes and trucks, which significantly reduces the cost, complexity and risk. *** In DD OS, data is deduplicated as it is written to the source system and replication preserves deduplication. This ensures that the network is efficiently utilized for creating a DR copy of backup and archive data.").

15. The Accused Instrumentalities further include "a controller configured to edit a list so that content data is registered in the list." For example, Data Domain Replicator includes both command-line and graphical tools to select a directory, managed file, or MTree to be replicated, and register. The tools register the selected content in a list. *See, e.g.*, Administration Guide ("You can manage replication using the Data Domain System Manager (DD System Manager) or the Data Domain Operating System (DD OS) Command Line Interface (CLI)."); H7082:

**MANAGED FILE REPLICATION**

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

***

**CONTENT AWARE REPLICATION**

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

***

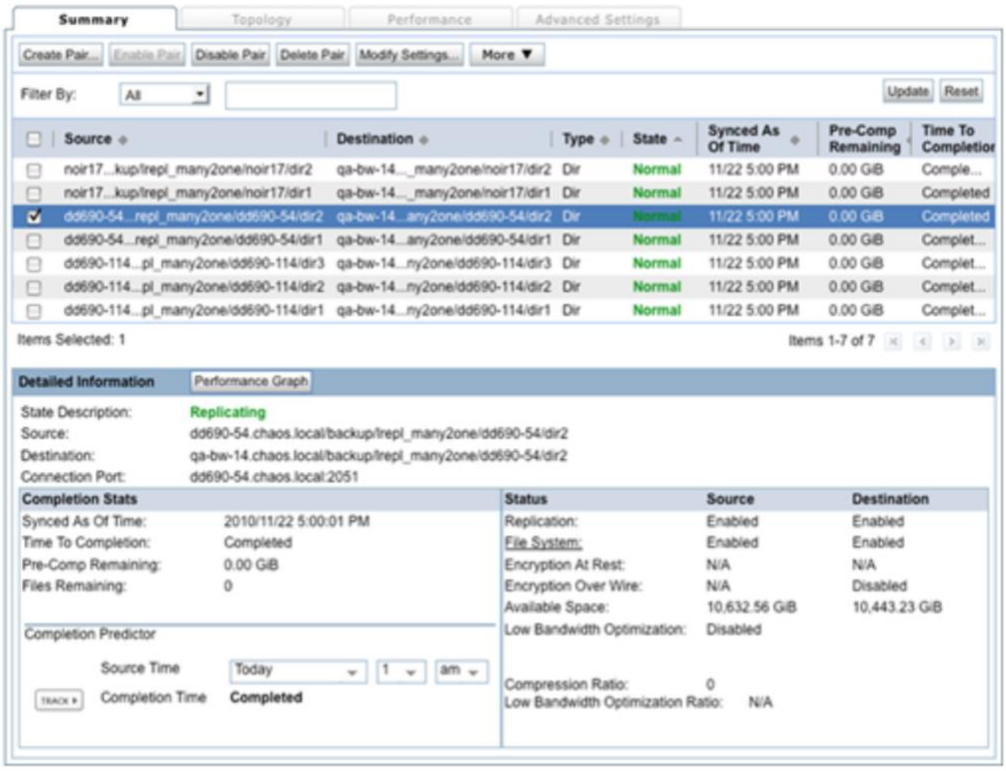


Figure 12: Detailed views of replication configuration and status

16. The Accused Instrumentalities further include a controller configured "to uniquely associate the list with the external apparatus using a unique identification of the external apparatus." For example, a particular unit of content data can be replicated to one or multiple destinations. Because it is possible to replicate some data to one destination, and different data to a different destination, there is a unique list associated with each destination. The destination is identified with a unique identifier, e.g. system name or host name. *See, e.g.*, Administration Guide:

### One-to-many replication

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

**Figure 23** One-to-many replication

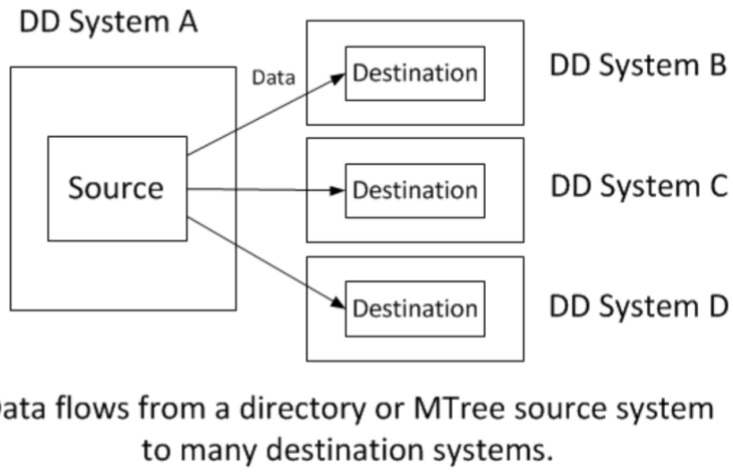


| | |
|---|---|
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |

*See also* H7082:

**FLEXIBLE REPLICATION TOPOLOGIES**

To enable enterprise-wide data protection, DD Replicator provides multiple replication topologies – system mirroring, selective data replication, bi-directional replication, many-to-one replication, one-to-many replication and cascaded replication (see Figure 10). With many-to-one replication, up to 540 Data Domain systems in geographically distributed locations can replicate into a single DD9800 at the central data center.

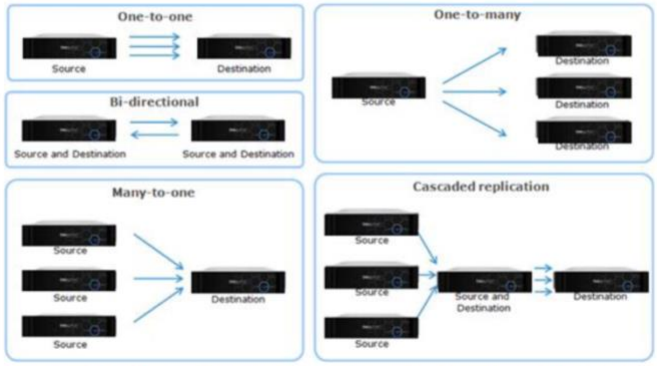


Figure 10: Replication supports a wide variety of topology DR needs

*See also* "Data Domain Advanced Features and Functions: Velocity Partner Accreditation," *available at* https://web.archive.org/web/20161125172552/

http://www.slideshare.net:80/solarisyougood/emc-data-domain-advanced-features-and-functions:

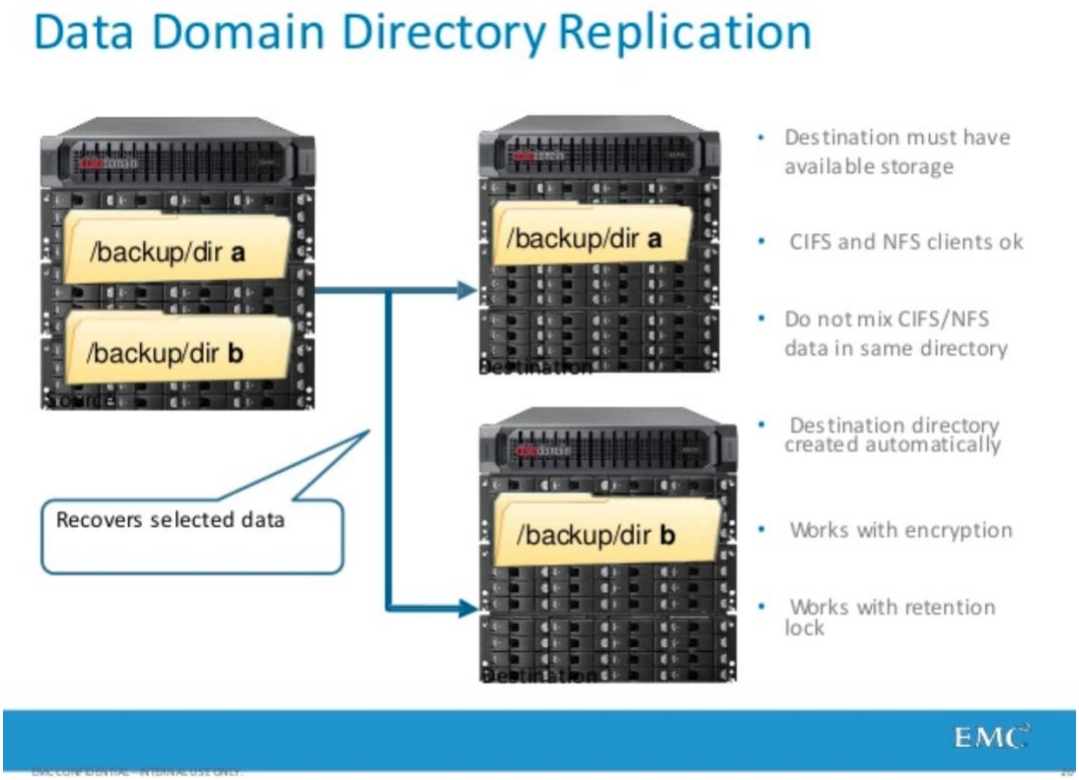


17. The Accused Instrumentalities further include a controller configured "to extract the list associated with the external apparatus from a plurality of lists in the communication apparatus when the external apparatus is connected to the communication apparatus." For example, when the destination system is connected to the source Data Domain Replicator system, the source system identifies the data to be replicated to the destination. Because it is possible to replicate some data to one destination, and different data to a different destination, there is a unique list associated with each destination; and there is a plurality of lists corresponding to the plurality of destinations. *See, e.g.*, Administration Guide:

> "To start replication between a source and destination, use the replication initialize command on the source. This command checks that the configuration and connections are correct and returns error messages if any problems appear."

**One-to-many replication**

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

**Figure 23** One-to-many replication

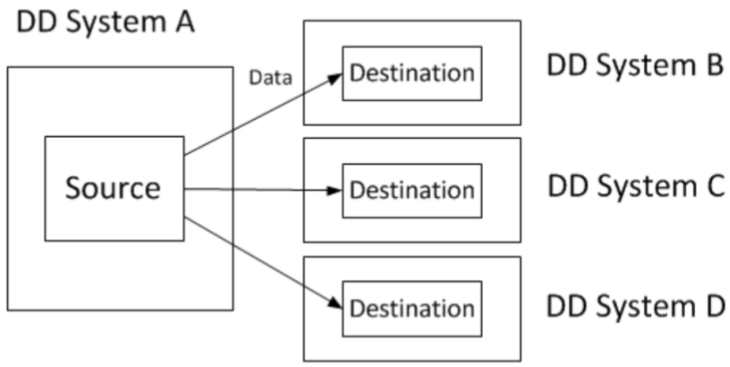


Data flows from a directory or MTree source system to many destination systems.

| | |
|---|---|
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |

*See also* H7082:

**FLEXIBLE REPLICATION TOPOLOGIES**

To enable enterprise-wide data protection, DD Replicator provides multiple replication topologies – system mirroring, selective data replication, bi-directional replication, many-to-one replication, one-to-many replication and cascaded replication (see Figure 10). With many-to-one replication, up to 540 Data Domain systems in geographically distributed locations can replicate into a single DD9800 at the central data center.

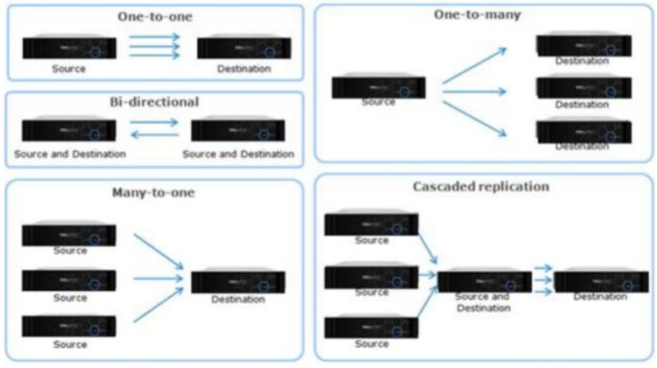


Figure 10: Replication supports a wide variety of topology DR needs

*See also* "Data Domain Advanced Features and Functions: Velocity Partner Accreditation":

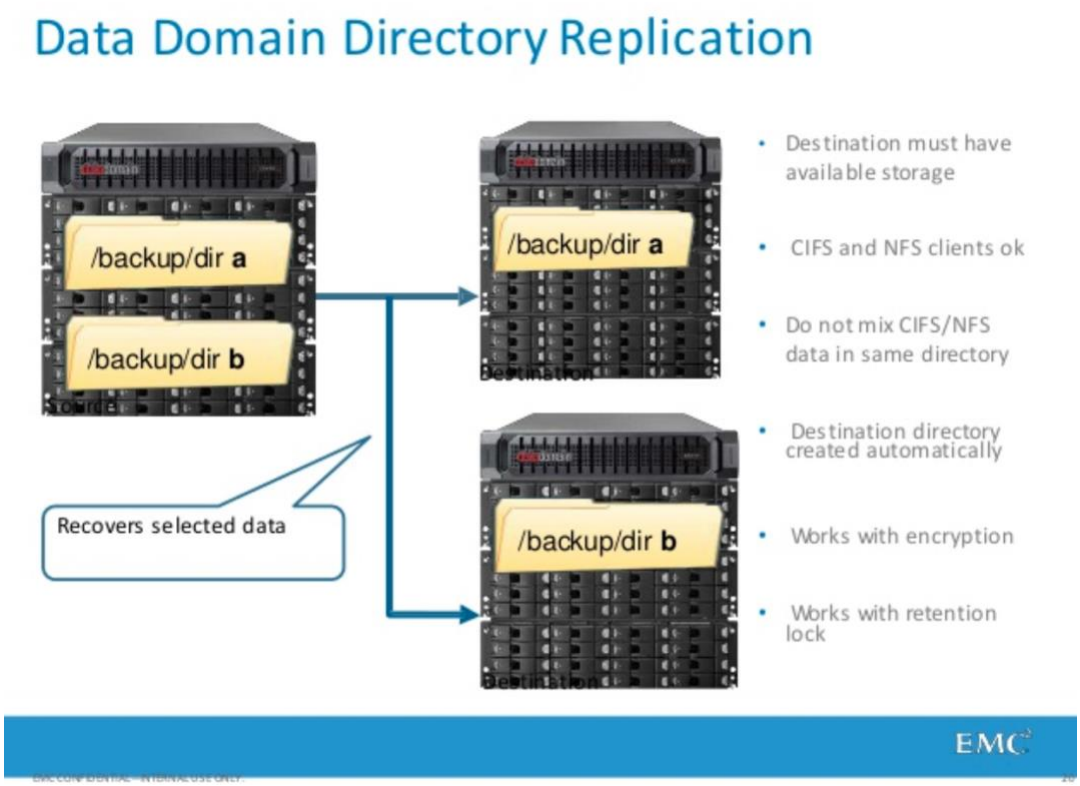


18. The Accused Instrumentalities further include a controller configured "to control transferring of content data registered in the extracted list to the external apparatus." For example, the Accused Instrumentalities control data replication to ensure that only the data on the list is replicated, and only when the first and second apparatuses are connected (e.g. not in an erroneous "Disconnected" state). *See, e.g.*, Administration Guide:

| Item | Description |
|---|---|
| Source | System and path name of the source context, with format *system.path*. For example, for directory `dir1` on `system dd120-22`, you would see `dd120-22.chaos.local/data/col1/dir1`. |
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |
| Type | Type of context: MTree, directory (Dir), or Pool. |
| State | Possible states of replication pair status include:<br>• Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating.<br>• Idle – For MTree replication, this state can display if the replication process is not currently active or for network errors (such as the destination system being inaccessible).<br>• Warning – If there is an unusual delay for the first five states, or for the Uninitialized state.<br>• Error – Any possible error states, such as Disconnected. |
| Synced As Of Time | Timestamp for last automatic replication sync operation performed by the source. For MTree replication, this value is updated when a snapshot is exposed on the destination. For directory replication, it is updated when a sync point inserted by the source is applied. A value of unknown displays during replication initialization. |

*See also* H7082:

**MANAGED FILE REPLICATION**

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

***

**CONTENT AWARE REPLICATION**

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

***

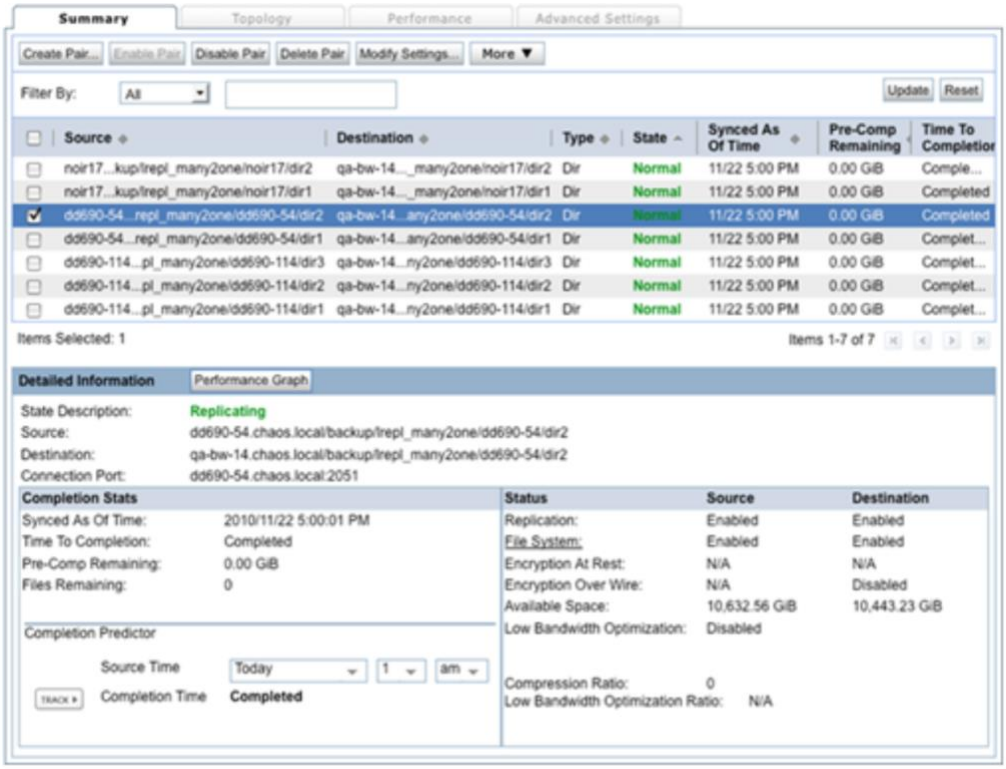


Figure 12: Detailed views of replication configuration and status

19. Each defendant has had knowledge of the ’581 Patent and its infringement since at least the filing of the original Complaint in this action, or shortly thereafter, including by way of this lawsuit. By the time of trial, each defendant will have known and intended (since receiving such notice) that its continued actions would actively induce and contribute to the infringement of the claims of the ‘’581 Patent.

20. Each defendant’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users

of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the claims of the '581 Patent. Use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the claims of the '581 Patent.

21. For example, each defendant explains to customers the benefits of using the Accused Instrumentalities, such as by touting their advantages of replicating data among multiple devices in the case of Data Domain, or of replicating data among multiple sites in the case of RecoverPoint. Each defendant also induces its customers to use the Accused Instrumentalities to infringe other claims of the '581 Patent. Each defendant specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the '581 Patent. Each defendant performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '581 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, each defendant engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.,* through its user manuals, product support, marketing materials, demonstrations, installation support, and training materials to actively induce the users of the accused products to infringe the '581 Patent. Accordingly, each defendant has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the '581 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the '581 Patent. Accordingly, each defendant has been (since at least as of filing of the original complaint), and currently is, inducing infringement of the '581 Patent, in violation of 35 U.S.C. § 271(b).

22. Each defendant has also infringed, and continues to infringe, claims of the '581 Patent by offering to commercially distribute, commercially distributing, making, and/or importing the Accused Instrumentalities, which are used in practicing the process,

or using the systems, of the ’581 Patent, and constitute a material part of the invention. Defendant knows the components in the Accused Instrumentalities to be especially made or especially adapted for use in infringement of the ’581 Patent, not a staple article, and not a commodity of commerce suitable for substantial noninfringing use. For example, the ordinary way of using the Accused Instrumentalities infringes the patent claims, and as such, is especially adapted for use in infringement. Accordingly, each defendant has been, and currently is, contributorily infringing the ’581 Patent, in violation of 35 U.S.C. § 271(c).

23. For similar reasons, each defendant also infringes the ’581 Patent by supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the Accused Instrumentalities, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the ’581 Patent if such combination occurred within the United States. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., Data Domain servers, RecoverPoint appliances) and software (e.g., Data Domain OS, RecoverPoint software) components of the Accused Instrumentalities in such a manner as to actively induce the combination of such components (e.g., by instructing users to combine multiple Data Domain or RecoverPoint servers into an infringing system) outside of the United States.

24. Each defendant also indirectly infringes the ’581 Patent by supplying or causing to be supplied in or from the United States components of the Accused Instrumentalities that are especially made or especially adapted for use in infringing the ’581 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use, and where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components are combined outside of the United States in a manner that would infringe the ’581 Patent if such combination occurred within the United States. Because the Accused

Instrumentalities are designed to operate as the claimed system and apparatus, the Accused Instrumentalities have no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., separate Data Domain servers, separate RecoverPoint appliances) and software (e.g., Data Domain OS, RecoverPoint software) components that are especially made or especially adapted for use in the Accused Instrumentalities, where such hardware and software components are not staple articles or commodities of commerce suitable for substantial noninfringing use, knowing that such components are so made or adapted and intending that such components are combined outside of the United States, as evidenced by each defendant's own actions or instructions to users in, e.g., combining multiple Data Domain or RecoverPoint servers into infringing systems, and enabling and configuring the infringing functionalities of the Accused Instrumentalities.

25. As a result of Defendant's infringement of the '581 Patent, Plaintiff Data Scape is entitled to monetary damages in an amount adequate to compensate for each Defendant's infringement, but in no event less than a reasonable royalty for the use made of the invention by each Defendant, together with interest and costs as fixed by the Court.

## COUNT II

## INFRINGEMENT OF U.S. PATENT NO. 7,720,929

26. Data Scape is the owner by assignment of United States Patent No. 7,720,929 ("the '929 Patent"), entitled "Communication System And Its Method and Communication Apparatus And Its Method." The '929 Patent was duly and legally issued by the United States Patent and Trademark Office on May 18, 2010. A true and correct copy of the '929 Patent is included as Exhibit B.

27. Each defendant has offered for sale, sold and/or imported into the United States products and services that infringe the '929 patent, and continues to do so. By way

of illustrative example, these infringing products and services include, without limitation, Defendant's products and services, *e.g.*, Data Domain software and hardware, RecoverPoint software and hardware, and all versions and variations thereof since the issuance of the '929 Patent ("Accused Instrumentalities").

28. Each defendant has directly infringed and continues to infringe the '929 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities. Each defendant uses the Accused Instrumentalities for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

29. For example, the Accused Instrumentalities infringe Claim 1 (as well as other claims) of the '929 Patent. One non-limiting example of the Accused Instrumentalities' infringement is presented below:

30. The Accused Instrumentalities include "[a] communication system including a first apparatus having a first storage medium, and a second apparatus." For example, the Accused Instrumentalities communicate data stored on one device (e.g. a Data Domain system) to another device (e.g. a second Data Domain system). *See, e.g.*, EMC Data Domain Operating System Version 5.7 Administration Guide ("Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of management. *** The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. A Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination for one or more pairs. After replication is started, the source system automatically sends any new backup data to the destination system.").

31. The Accused Instrumentalities include a second apparatus comprising: "a second storage medium configured to store management information of data to be transferred to said first storage medium." For example, a Data Domain system includes disks and/or solid-state storage medium. *See, e.g.*, Administration Guide ("Storage in most Data Domain systems is set up in a double parity RAID 6 configuration (two parity drives). Additionally, most configurations include a hot spare in each enclosure, except the DD1xx series systems, which use eight disks. *** To keep data synchronized during a hardware or power failure, the Data Domain system uses NVRAM (non-volatile RAM) to track outstanding I/O operations. An NVRAM card with fully charged batteries (the typical state) can retain data for a period of hours, which is determined by the hardware in use."). The second storage medium is configured to store management information of data to be transferred, e.g. replication configuration settings, folder metadata, etc.

32. The Accused Instrumentalities further include a second apparatus comprising "a communicator configured to communicate with said first apparatus." For example, a Data Domain system can connect to another Data Domain system over a wide area network. *See, e.g.*, Whitepaper H7082: Dell EMC Data Domain Replicator ("In comparison, replication uses the wide area network (WAN) as the transport mechanism for data instead of tapes and trucks, which significantly reduces the cost, complexity and risk. *** In DD OS, data is deduplicated as it is written to the source system and replication preserves deduplication. This ensures that the network is efficiently utilized for creating a DR copy of backup and archive data.").

33. The Accused Instrumentalities further include a second apparatus comprising "a detector configured to detect whether said first apparatus and a second apparatus are connected." For example, Data Domain Replicator uses a detector to determine whether it is connected to the first apparatus. *See, e.g.*, Administration Guide ("To start replication between a source and destination, use the replication initialize

command on the source. This command checks that the configuration and connections are correct and returns error messages if any problems appear.").

34. The Accused Instrumentalities further include a second apparatus comprising "an editor configured to select certain data to be transferred and to edit said management information based on said selection without regard to the connection of said first apparatus." For example, Data Domain Replicator includes both command-line and graphical tools to select a directory, managed file, or MTree to be transferred. The tools edit internal configuration information to maintain those settings. On information and belief, the editing occurs without regard to the connection of the first apparatus. *See, e.g.*, Administration Guide ("You can manage replication using the Data Domain System Manager (DD System Manager) or the Data Domain Operating System (DD OS) Command Line Interface (CLI)."); H7082:

**MANAGED FILE REPLICATION**

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

***

**CONTENT AWARE REPLICATION**

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

***

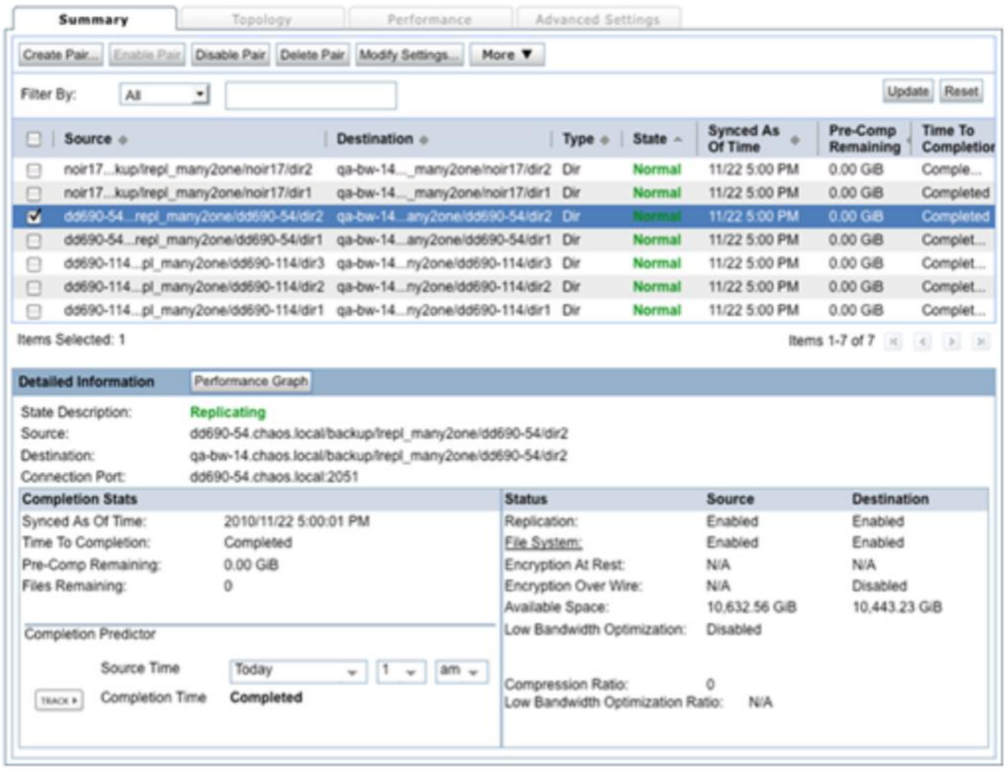


Figure 12: Detailed views of replication configuration and status

35. The Accused Instrumentalities further include a second apparatus comprising "a controller configured to control transfer of the selected data stored in said second apparatus to said first apparatus via said communicator based on said management information edited by said editor when said detector detects that said first apparatus and said second apparatus are connected." For example, the Accused Instrumentalities control data replication to ensure that only the data selected by the editor is replicated, and only when the first and second apparatuses are connected (e.g. not in an erroneous "Disconnected" state). *See, e.g.*, Administration Guide:

| Item | Description |
|---|---|
| Source | System and path name of the source context, with format *system.path*. For example, for directory `dir1` on `system dd120-22`, you would see `dd120-22.chaos.local/data/col1/dir1`. |
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |
| Type | Type of context: MTree, directory (Dir), or Pool. |
| State | Possible states of replication pair status include:<br>• Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating.<br>• Idle – For MTree replication, this state can display if the replication process is not currently active or for network errors (such as the destination system being inaccessible).<br>• Warning – If there is an unusual delay for the first five states, or for the Uninitialized state.<br>• Error – Any possible error states, such as Disconnected. |
| Synced As Of Time | Timestamp for last automatic replication sync operation performed by the source. For MTree replication, this value is updated when a snapshot is exposed on the destination. For directory replication, it is updated when a sync point inserted by the source is applied. A value of unknown displays during replication initialization. |

*See also* H7082:

**MANAGED FILE REPLICATION**

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

***

**CONTENT AWARE REPLICATION**

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

***

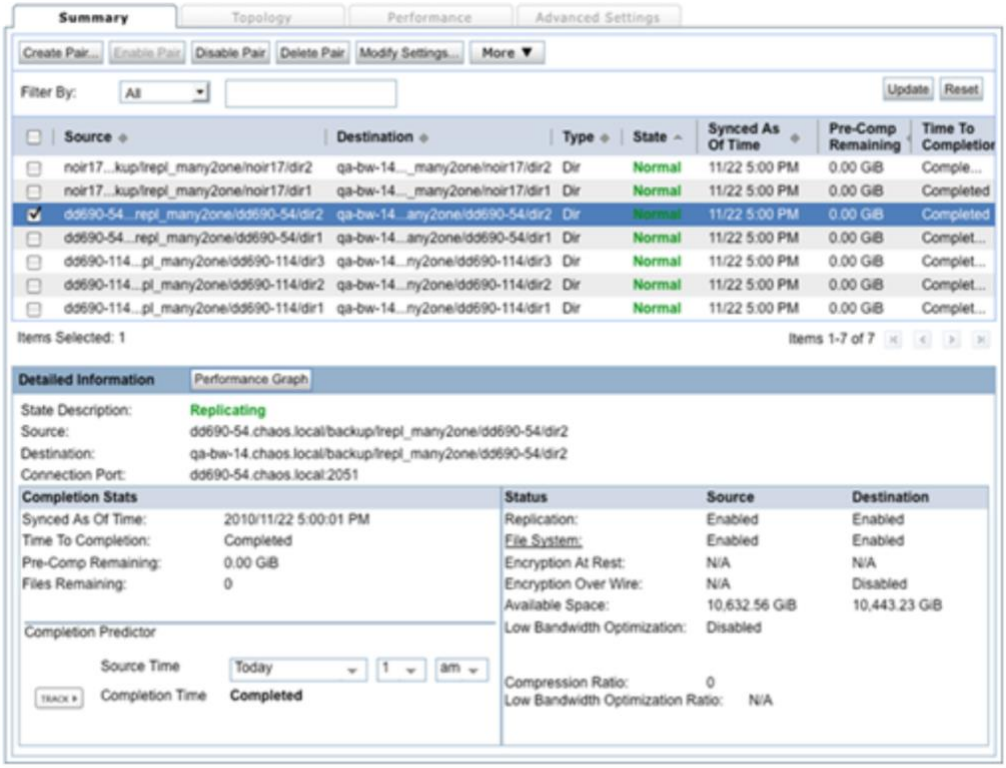


Figure 12: Detailed views of replication configuration and status

36. The Accused Instrumentalities further include a second apparatus "wherein said controller is configured to compare said management information edited by said editor with management information of data stored in said first storage medium and to transmit data in said second apparatus based on result of the comparison." For example, the "DD source system" replicates, i.e. transmits, data from the source system to the destination system based on management information of data stored in the first system. *See, e.g.*, H7082:

The simplest type of replication is from a DD source system to a DD destination system, otherwise known as a *one-to-one* replication pair. This replication topology can be configured with directory, MTree, or collection replication types.

**Figure 21** One-to-one replication pair

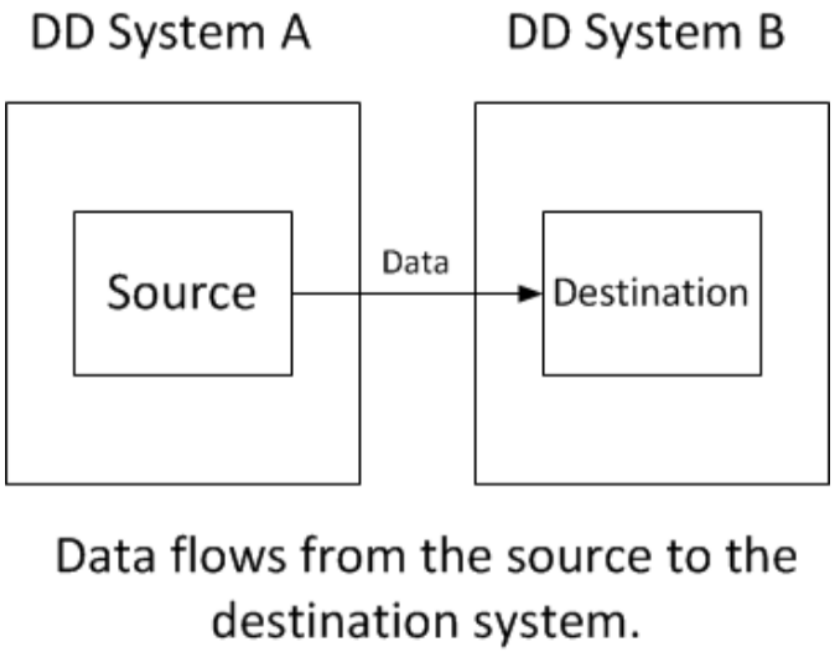


***

In Figure 3 below, Metadata exchange between the source and destination ensures that a data segment only needs to be sent to the destination once, irrespective of where the data comes from. This provides significant efficiencies over the WAN in many-to-one deployments since common segments on different sources only need to be sent once.

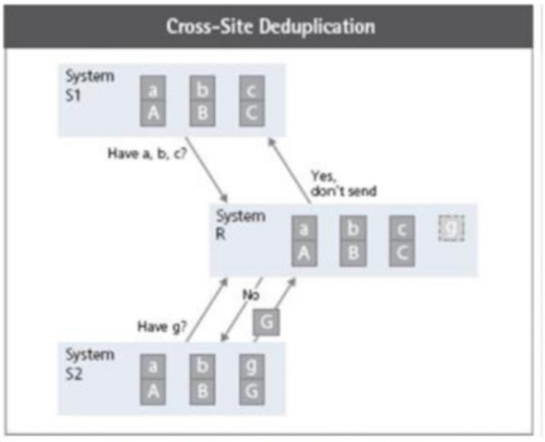


Figure 3: Cross-Site Deduplication

37. Each defendant has had knowledge of the '929 Patent and its infringement since at least the filing of the original Complaint in this action, or shortly thereafter, including by way of this lawsuit. By the time of trial, each defendant will have known and intended (since receiving such notice) that its continued actions would actively induce and contribute to the infringement of the claims of the '929 Patent.

38. Each defendant's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the claims of the '929 Patent. Use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the claims of the '929 Patent.

39. For example, each defendant explains to customers the benefits of using the Accused Instrumentalities, such as by touting their advantages of replicating data among multiple devices in the case of Data Domain, or of replicating data among multiple sites in the case of RecoverPoint. Each defendant also induces its customers to use the Accused Instrumentalities to infringe other claims of the '929 Patent. Each defendant specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the '929 Patent. Each defendant performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '929 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, each defendant engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through its user manuals, product support, marketing materials, demonstrations, installation support, and training materials to actively induce the users of the accused products to infringe the '929 Patent. Accordingly, each defendant has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems

infringing the ’929 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the ’929 Patent. Accordingly, each defendant has been (since at least as of filing of the original complaint), and currently is, inducing infringement of the ’929 Patent, in violation of 35 U.S.C. § 271(b).

40. Each defendant has also infringed, and continues to infringe, claims of the ’929 Patent by offering to commercially distribute, commercially distributing, making, and/or importing the Accused Instrumentalities, which are used in practicing the process, or using the systems, of the ’929 Patent, and constitute a material part of the invention. Defendant knows the components in the Accused Instrumentalities to be especially made or especially adapted for use in infringement of the ’929 Patent, not a staple article, and not a commodity of commerce suitable for substantial noninfringing use. For example, the ordinary way of using the Accused Instrumentalities infringes the patent claims, and as such, is especially adapted for use in infringement. Accordingly, each defendant has been, and currently is, contributorily infringing the ’929 Patent, in violation of 35 U.S.C. § 271(c).

41. For similar reasons, each defendant also infringes the ’929 Patent by supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the Accused Instrumentalities, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the ’929 Patent if such combination occurred within the United States. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., Data Domain servers, RecoverPoint appliances) and software (e.g., Data Domain OS, RecoverPoint software) components of the Accused Instrumentalities in such a manner as to actively induce the combination of such components (e.g., by instructing users to combine multiple Data Domain or RecoverPoint servers into an infringing system) outside of the United States.

42. Each defendant also indirectly infringes the ’929 Patent by supplying or causing to be supplied in or from the United States components of the Accused Instrumentalities that are especially made or especially adapted for use in infringing the ’929 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use, and where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components are combined outside of the United States in a manner that would infringe the ’929 Patent if such combination occurred within the United States. Because the Accused Instrumentalities are designed to operate as the claimed system and apparatus, the Accused Instrumentalities have no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., separate Data Domain servers, separate RecoverPoint appliances) and software (e.g., Data Domain OS, RecoverPoint software) components that are especially made or especially adapted for use in the Accused Instrumentalities, where such hardware and software components are not staple articles or commodities of commerce suitable for substantial noninfringing use, knowing that such components are so made or adapted and intending that such components are combined outside of the United States, as evidenced by each defendant’s own actions or instructions to users in, e.g., combining multiple Data Domain or RecoverPoint servers into infringing systems, and enabling and configuring the infringing functionalities of the Accused Instrumentalities.

43. As a result of Defendant’s infringement of the ’929 Patent, Plaintiff Data Scape is entitled to monetary damages in an amount adequate to compensate for each Defendant’s infringement, but in no event less than a reasonable royalty for the use made of the invention by each Defendant, together with interest and costs as fixed by the Court.

## COUNT III

## INFRINGEMENT OF U.S. PATENT NO. 7,617,537

44. Data Scape is the owner by assignment of United States Patent No. 7,617,537 ("the '537 Patent"), entitled "Communication System And Its Method and Communication Apparatus And Its Method." The '537 Patent was duly and legally issued by the United States Patent and Trademark Office on November 10, 2009. A true and correct copy of the '537 Patent is included as Exhibit C.

45. Each defendant has offered for sale, sold and/or imported into the United States products and services that infringe the '537 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, Defendant's products and services, *e.g.*, Data Domain software and hardware, and all versions and variations thereof since the issuance of the '537 Patent ("Accused Instrumentalities").

46. Each defendant has directly infringed and continues to infringe the '537 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities. Each defendant uses the Accused Instrumentalities for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

47. For example, the Accused Instrumentalities infringe Claim 43 (as well as other claims) of the '537 Patent. One non-limiting example of the Accused Instrumentalities' infringement is presented below:

48. The Accused Instrumentalities include "[a] computer readable storage medium encoded with computer program instructions executable by a computer to implement a method of transferring content data to a first apparatus from a second apparatus." For example, the Accused Instrumentalities communicate data stored on one device (e.g. a Data Domain system) to another device (e.g. a second Data Domain system).

*See, e.g.*, EMC Data Domain Operating System Version 5.7 Administration Guide ("Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of management. *** The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. A Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination for one or more pairs. After replication is started, the source system automatically sends any new backup data to the destination system.").

49. The Accused Instrumentalities include instructions that "judge whether said first apparatus and said second apparatus are connected." For example, Data Domain Replicator determines whether it is connected to the first apparatus. *See, e.g.*, Administration Guide ("To start replication between a source and destination, use the replication initialize command on the source. This command checks that the configuration and connections are correct and returns error messages if any problems appear.").

50. The Accused Instrumentalities include instructions that "compare, upon judging that said first apparatus and said second apparatus are connected, an identifier of said first apparatus with a corresponding identifier in said second apparatus." For example, the destination system is identified with a unique identifier, e.g. system name or host name. *See, e.g.*, Administration Guide:

| | |
|---|---|
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |

51. The Accused Instrumentalities include instructions that "compare, when said identifier of said first apparatus corresponds to said identifier stored in said second apparatus, a first list of content data of said first apparatus and a second list of content data of said second apparatus." For example, the source system identifies the replication sets

associated with the destination system, using the unique identifier, e.g. system name or host name, associated with the destination system. The source system then identifies a list of content data in the destination system and compares the lists to determine what content to send, e.g. by deduplicating the data or by only transmitting changes to the data. *See, e.g.*, Whitepaper H7082: Dell EMC Data Domain Replicator:

> "With directory replication, a replication context pairs a directory (and all files and directories below it) on a source system with a destination directory on a different system, as seen in Figure 2. During replication, deduplication is preserved since data segments that already reside on the destination system will not be resent across the WAN."
>
> "Metadata exchange between the source and destination ensures that a data segment only needs to be sent to the destination once, irrespective of where the data comes from. This provides significant efficiencies over the WAN in many-to-one deployments since common segments on different sources only need to be sent once."
>
> "MTree replication has a lot of commonality with directory replication. It uses the same WAN deduplication mechanism used by directory replication to avoid sending redundant data over the WAN. It also supports all the topologies supported by directory replication (one-to-one, bi-directional, one-to-many, many-to-one, cascaded). In addition, one can configure MTree replication to replicate MTree data on a system that already leverages directory replication and/or managed file replication."

*See also* Administration Guide:

**One-to-many replication**

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

**Figure 23** One-to-many replication

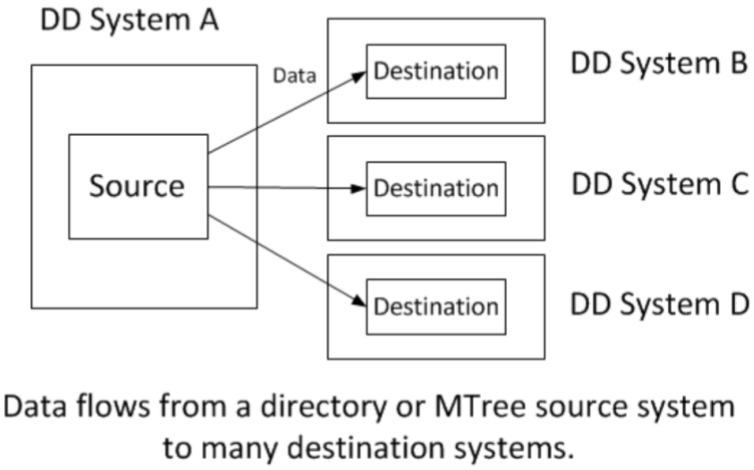


> "With cross-site deduplication, any redundant segment previously transferred by any other site, or as a result of a local backup or archive, will not be replicated again. This improves network efficiency across all sites and reduces daily network bandwidth requirements up to 99%, making network-based replication fast, reliable, and cost- effective."

52. The Accused Instrumentalities include instructions to "transfer first content data, from the second apparatus to the first apparatus, which is registered in said second list and is not registered in said first list." For example, the source system identifies the replication sets associated with the destination system, using the unique identifier, e.g. system name or host name, associated with the destination system. The source system then identifies a list of content data in the destination system and compares the lists to determine what content to send, e.g. by deduplicating the data or by only transmitting changes to the data. *See, e.g.*, Whitepaper H7082: Dell EMC Data Domain Replicator:

> "With directory replication, a replication context pairs a directory (and all files and directories below it) on a source system with a destination directory on a different system, as seen in Figure 2. During replication, deduplication is preserved since data segments that already reside on the destination system will not be resent across the WAN."

"Metadata exchange between the source and destination ensures that a data segment only needs to be sent to the destination once, irrespective of where the data comes from. This provides significant efficiencies over the WAN in many-to-one deployments since common segments on different sources only need to be sent once."

"MTree replication has a lot of commonality with directory replication. It uses the same WAN deduplication mechanism used by directory replication to avoid sending redundant data over the WAN. It also supports all the topologies supported by directory replication (one-to-one, bi-directional, one-to-many, many-to-one, cascaded). In addition, one can configure MTree replication to replicate MTree data on a system that already leverages directory replication and/or managed file replication."

*See also* Administration Guide:

**One-to-many replication**

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

**Figure 23** One-to-many replication

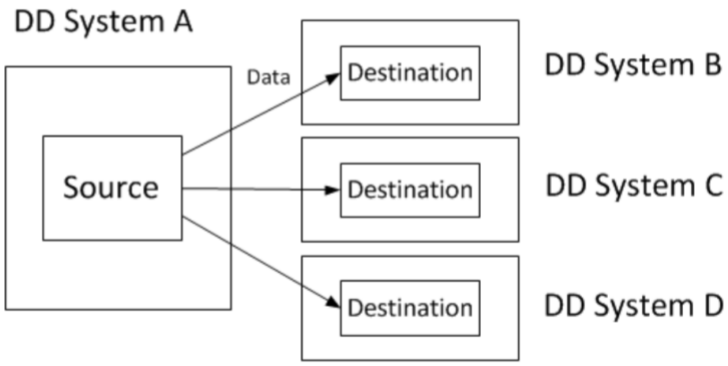


Data flows from a directory or MTree source system to many destination systems.

"With cross-site deduplication, any redundant segment previously transferred by any other site, or as a result of a local backup or archive, will not be replicated again. This improves network efficiency across all sites and reduces daily network bandwidth requirements up to 99%, making network-based replication fast, reliable, and cost- effective."

53. The Accused Instrumentalities include instructions to "delete second content data, from the first apparatus, which is registered in said first list and is not registered in said second list." For example, when a file or snapshot is deleted from the

source system, the DD Replicator system will replicate that deletion to the destination system, deleting the associated content data from the destination system. *See, e.g.*, "How MTree Replication Works," *available at* https://community.emc.com/docs/DOC-39126

> "One of the fundamental differences between MTree and Directory replication is the method used for determining what needs to be replicated from the source to the destination. MTree replication leverages snapshots to ensure that the destination Data Domain system will always be a point-in-time image of the source Data Domain system. Snapshots are point-in-time images of the MTree. For each MTree replication context the system will auto-create (and auto-delete) snapshots of the MTree.
>
> In contrast, Directory replication does not use snapshots, it uses a replication log. It must replicate every change that has been done to the content of the source directory in the sequence that it was changed. So for example, if file_1 is created, then modified, and then deleted. The directory replication will replicate the creation of file_1, the modification, and the deletion steps to the destination DDR."

54. Each defendant has had knowledge of the '537 Patent and its infringement since at least the filing of the original Complaint in this action, or shortly thereafter, including by way of this lawsuit. By the time of trial, each defendant will have known and intended (since receiving such notice) that its continued actions would actively induce and contribute to the infringement of the claims of the '537 Patent.

55. Each defendant's affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the claims of the '537 Patent. Use of the Accused Instrumentalities in their ordinary and customary fashion results in infringement of the claims of the '537 Patent.

56. For example, each defendant explains to customers the benefits of using the Accused Instrumentalities, such as by touting their advantages of replicating data among multiple devices. Each defendant also induces its customers to use the Accused Instrumentalities to infringe other claims of the '537 Patent. Each defendant specifically

intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the '537 Patent. Each defendant performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the '537 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, each defendant engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through its user manuals, product support, marketing materials, demonstrations, installation support, and training materials to actively induce the users of the accused products to infringe the '537 Patent. Accordingly, each defendant has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the '537 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the '537 Patent. Accordingly, each defendant has been (since at least as of filing of the original complaint), and currently is, inducing infringement of the '537 Patent, in violation of 35 U.S.C. § 271(b).

57. Each defendant has also infringed, and continues to infringe, claims of the '537 Patent by offering to commercially distribute, commercially distributing, making, and/or importing the Accused Instrumentalities, which are used in practicing the process, or using the systems, of the '537 Patent, and constitute a material part of the invention. Defendant knows the components in the Accused Instrumentalities to be especially made or especially adapted for use in infringement of the '537 Patent, not a staple article, and not a commodity of commerce suitable for substantial noninfringing use. For example, the ordinary way of using the Accused Instrumentalities infringes the patent claims, and as such, is especially adapted for use in infringement. Accordingly, each defendant has been, and currently is, contributorily infringing the '537 Patent, in violation of 35 U.S.C. § 271(c).

58. For similar reasons, each defendant also infringes the '537 Patent by supplying or causing to be supplied in or from the United States all or a substantial portion

of the components of the Accused Instrumentalities, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the ’537 Patent if such combination occurred within the United States. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., Data Domain servers) and software (e.g., Data Domain OS) components of the Accused Instrumentalities in such a manner as to actively induce the combination of such components (e.g., by instructing users to combine multiple Data Domain servers into an infringing system) outside of the United States.

59. Each defendant also indirectly infringes the ’537 Patent by supplying or causing to be supplied in or from the United States components of the Accused Instrumentalities that are especially made or especially adapted for use in infringing the ’537 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use, and where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components are combined outside of the United States in a manner that would infringe the ’537 Patent if such combination occurred within the United States. Because the Accused Instrumentalities are designed to operate as the claimed system and apparatus, the Accused Instrumentalities have no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., separate Data Domain servers) and software (e.g., Data Domain OS) components that are especially made or especially adapted for use in the Accused Instrumentalities, where such hardware and software components are not staple articles or commodities of commerce suitable for substantial noninfringing use, knowing that such components are so made or adapted and intending that such components are combined outside of the United States, as evidenced by each defendant’s own actions

or instructions to users in, e.g., combining multiple Data Domain servers into infringing systems, and enabling and configuring the infringing functionalities of the Accused Instrumentalities.

60. As a result of Defendant's infringement of the '537 Patent, Plaintiff Data Scape is entitled to monetary damages in an amount adequate to compensate for each Defendant's infringement, but in no event less than a reasonable royalty for the use made of the invention by each Defendant, together with interest and costs as fixed by the Court.

**COUNT IV**

**INFRINGEMENT OF U.S. PATENT NO. 9,715,893**

61. Data Scape is the owner by assignment of United States Patent No. 9,715,893 ("the '893 Patent"), entitled "Recording Apparatus, Server Apparatus, Recording Method, Program and Storage Medium." The '893 Patent was duly and legally issued by the United States Patent and Trademark Office on July 25, 2017. A true and correct copy of the '893 Patent is included as Exhibit D.

62. Each defendant has offered for sale, sold and/or imported into the United States products and services that infringe the '893 patent, and continues to do so. By way of illustrative example, these infringing products and services include, without limitation, Defendant's products and services, *e.g.*, Data Domain software and hardware, and all versions and variations thereof since the issuance of the '893 Patent ("Accused Instrumentalities").

63. Each defendant has directly infringed and continues to infringe the '893 Patent, for example, by making, selling, offering for sale, and/or importing the Accused Instrumentalities, and through its own use and testing of the Accused Instrumentalities. Each defendant uses the Accused Instrumentalities for its own internal non-testing business purposes, while testing the Accused Instrumentalities, and while providing technical support and repair services for the Accused Instrumentalities to its customers.

64. For example, the Accused Instrumentalities infringe Claim 1 (as well as other claims) of the '893 Patent. One non-limiting example of the Accused Instrumentalities' infringement is presented below:

65. The Accused Instrumentalities include "[a] non-transitory computer-readable storage medium storing instructions which, when executed by a computer, cause the computer to perform a method of an information processing apparatus for transferring data." For example, the Accused Instrumentalities include instructions to transfer data stored on one device (e.g. a Data Domain system) to another device (e.g. a second Data Domain system). *See, e.g.,* EMC Data Domain Operating System Version 5.7 Administration Guide ("Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of management. *** The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. A Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination for one or more pairs. After replication is started, the source system automatically sends any new backup data to the destination system.").

66. The Accused Instrumentalities include instructions for "automatically reading first management data from a first storage medium, the first management data identifying files of source data stored on the first storage medium." For example, Data Domain Replicator includes management data, configured using either graphical or console commands, identifying files to be transferred. *See, e.g.,* Administration Guide ("You can manage replication using the Data Domain System Manager (DD System Manager) or the Data Domain Operating System (DD OS) Command Line Interface (CLI)."); H7082:

**DIRECTORY REPLICATION**

With directory replication, a replication context pairs a directory (and all files and directories below it) on a source system with a destination directory on a different system, as seen in Figure 2. During replication, deduplication is preserved since

data segments that already reside on the destination system will not be resent across the WAN. The destination directory will be read-only as long as the replication context is configured.

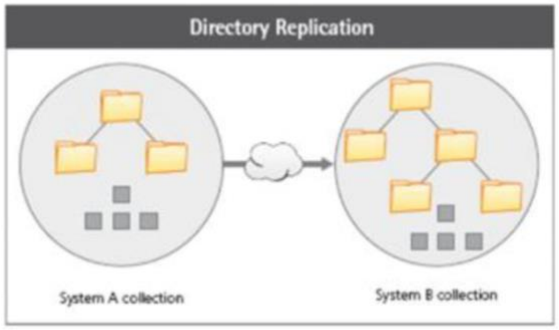


Figure 2: Data Domain Directory Replication

***

**MANAGED FILE REPLICATION**

Managed file replication using DD Boost allows the backup software to control the replication on a per-file basis. When integrated with DD Boost, the backup software's users can configure policies to selectively replicate the individual backup image or dataset to another system after completion of the backup. Unlike traditional vaulting or cloning to tape, the data is not read by the backup server to be written elsewhere. Instead, the backup software delegates the data movement to the DD system; thereby leveraging the most efficient method available to create a DR copy of the data.

The backup software decides when to get started, and knows when it is finished, based on interactive signaling between DD Boost and the Data Domain system. Using this approach, the backup software knows that the destination holds a copy of the file that is separate and different from the source's file, and retention periods for the two can be managed independently, for example, to keep full backups longer on the DR site. Furthermore, the backup operator has the flexibility to decide which backup images need to be replicated, and which ones do not require DR protection; e.g. user may decide that daily incremental backups do not need to be replicated, but weekly full backups should be replicated offsite.

***

**CONTENT AWARE REPLICATION**

Backup applications can write virtual synthetic full backups to Data Domain systems using the DD Boost protocol. The virtual full backups are synthesized from existing backups on the DD system, and provide significant performance improvements and network utilization reduction when writing backups.

DD Replicator applies the same synthesis optimizations for synthetic full backups to deliver similar performance improvements and network utilization reduction. The synthetic replication optimization is applicable with both Managed File replication and MTree replication.

67. The Accused Instrumentalities include instructions for "automatically identifying, by the computer, one of the files of source data based on the first management

data and second management data, the second management data identifying files of transferred data stored on a second storage medium, the one of the files of source data being absent from the second storage medium." For example, DD Replicator on the source system identifies files of source data that are absent from the destination system using management information, *e.g.* information associated with the deduplication or incremental transfer functionality of DD Replicator. *See, e.g.,* Whitepaper H7082: Dell EMC Data Domain Replicator:

> "With directory replication, a replication context pairs a directory (and all files and directories below it) on a source system with a destination directory on a different system, as seen in Figure 2. During replication, deduplication is preserved since data segments that already reside on the destination system will not be resent across the WAN."
>
> "Metadata exchange between the source and destination ensures that a data segment only needs to be sent to the destination once, irrespective of where the data comes from. This provides significant efficiencies over the WAN in many-to-one deployments since common segments on different sources only need to be sent once."
>
> "MTree replication has a lot of commonality with directory replication. It uses the same WAN deduplication mechanism used by directory replication to avoid sending redundant data over the WAN. It also supports all the topologies supported by directory replication (one-to-one, bi-directional, one-to-many, many-to-one, cascaded). In addition, one can configure MTree replication to replicate MTree data on a system that already leverages directory replication and/or managed file replication."

*See also* Administration Guide:

**One-to-many replication**

In one-to-many replication, data flows from a source directory or MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

**Figure 23** One-to-many replication

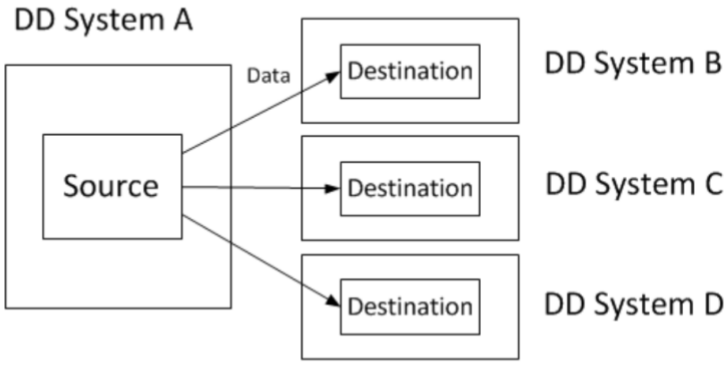


Data flows from a directory or MTree source system to many destination systems.

"With cross-site deduplication, any redundant segment previously transferred by any other site, or as a result of a local backup or archive, will not be replicated again. This improves network efficiency across all sites and reduces daily network bandwidth requirements up to 99%, making network-based replication fast, reliable, and cost- effective."

68. The Accused Instrumentalities include instructions for "automatically transferring the one of the files of source data to the second storage medium, the one of the files of source data being transferred becoming one of the files of transferred data." For example, the DD Replicator source system automatically transfers files to the destination system to ensure that the destination system contains a full replica. *See, e.g.*, EMC Data Domain Operating System Version 5.7 Administration Guide ("Data Domain system features ensure data integrity, reliable restoration, efficient resource usage, and ease of management. *** The EMC Data Domain Replicator sets up and manages the replication of backup data between two Data Domain systems. A Replicator pair consists of a source and a destination system and replicates a complete data set or directory from the source system to the destination system. An individual Data Domain system can be a part of multiple replication pairs and can serve as a source for one or more pairs and a destination

for one or more pairs. After replication is started, the source system automatically sends any new backup data to the destination system.").

69. The Accused Instrumentalities include instructions for "automatically displaying transferring status of the one of the files of source data by a symbolic figure." For example, the DD Replicator graphical user interface automatically displays file transfer status. *See, e.g.*, H7082:

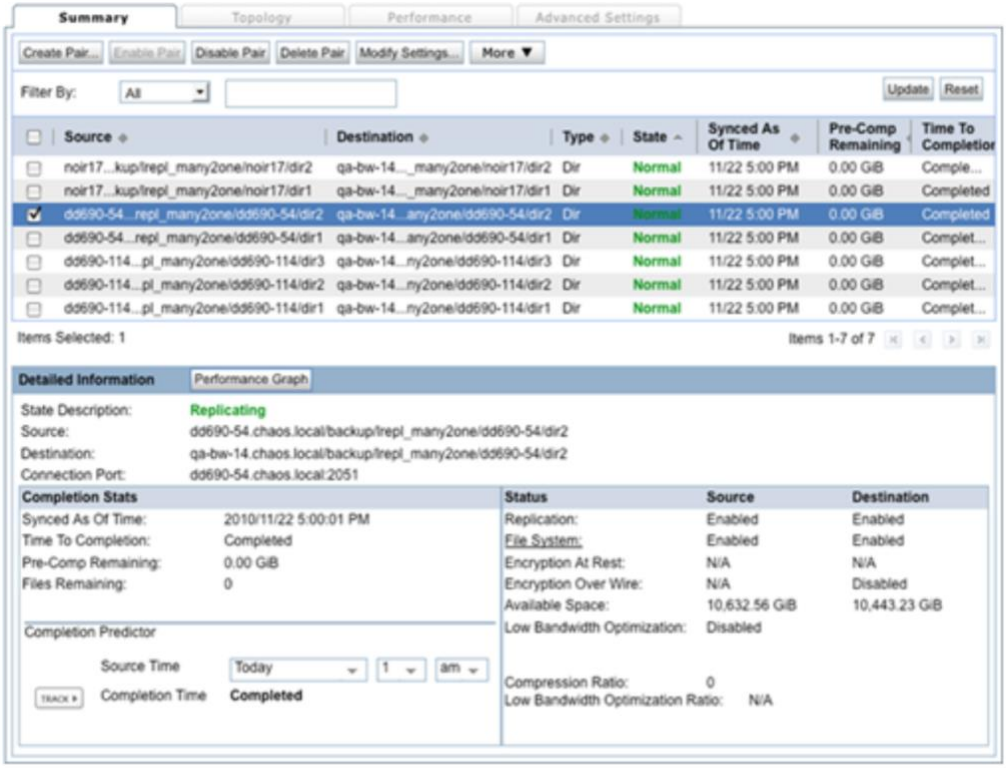


Figure 12: Detailed views of replication configuration and status

*See also, e.g.*, Administration Guide:

## Viewing estimated completion time for backup jobs

You can use the Completion Predictor to see the estimated time for when a backup replication job will be completed.

**Procedure**

1. Select **Replication** › **Summary**.
2. Select a Replication context for which to display Detailed Information.
3. In the Completion Predictor area, select options from the **Source Time** drop-down list for a replication's completion time, and select **Track**.

   The estimated time displays, in the Completion Time area, for when a particular backup job will finish its replication to the destination. If the replication is finished, the area shows `Completed`.

| Item | Description |
|---|---|
| Source | System and path name of the source context, with format *system.path*. For example, for directory `dir1` on `system dd120-22`, you would see `dd120-22.chaos.local/data/col1/dir1`. |
| Destination | System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd120-44`, you would see `dd120-44.chaos.local/data/col1/MTree1`. |
| Type | Type of context: MTree, directory (Dir), or Pool. |
| State | Possible states of replication pair status include:<br>• Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating.<br>• Idle – For MTree replication, this state can display if the replication process is not currently active or for network errors (such as the destination system being inaccessible).<br>• Warning – If there is an unusual delay for the first five states, or for the Uninitialized state.<br>• Error – Any possible error states, such as Disconnected. |
| Synced As Of Time | Timestamp for last automatic replication sync operation performed by the source. For MTree replication, this value is updated when a snapshot is exposed on the destination. For directory replication, it is updated when a sync point inserted by the source is applied. A value of unknown displays during replication initialization. |

Remote system files

The Show Details button provides information for the selected remote file replication system. File Replications shows starting and ending information, as well as size and data amount, for the selected remote file replication system. The Performance Graph shows performance over time for the selected remote file replication system.

**Table 171** File Replications

| Item | Description |
|---|---|
| Start | Starting point of time period. |
| End | Ending point of time period. |
| File Name | Name of specific replication file. |
| Status | Most recent status (Success, Failure). |
| Pre-Comp Size (MiB) | Amount of pre-compressed outbound and inbound data, as compared to network throughput or post-compressed data (in MiB). |
| Network Bytes (MiB) | Amount of network throughput data (in MiB). |

**Table 172** Performance Graph

| Item | Description |
|---|---|
| Duration | Duration for replication (either 1d, 7d or 30d). |
| Interval | Interval for replication (either Daily or Weekly). |
| Pre-Comp Replicated | Amount of pre-compressed outbound and inbound data (in GiB). |
| Post-Comp Replicated | Amount of post-compressed data (in GiB). |
| Network Bytes | Amount of network throughput data (in GiB). |
| Files Succeeded | Number of files that were successfully replicated. |
| Files Failed | Number of files that failed to be replicated. |
| Show in new window | Brings up a separate window. |
| Print | Prints the graph. |

70. Each defendant has had knowledge of the ’893 Patent and its infringement since at least the filing of the original Complaint in this action, or shortly thereafter, including by way of this lawsuit. By the time of trial, each defendant will have known and intended (since receiving such notice) that its continued actions would actively induce and contribute to the infringement of the claims of the ’893 Patent.

71. Each defendant’s affirmative acts of making, using, selling, offering for sale, and/or importing the Accused Instrumentalities have induced and continue to induce users of the Accused Instrumentalities to use the Accused Instrumentalities in their normal and customary way to infringe the claims of the ’893 Patent. Use of the Accused

Instrumentalities in their ordinary and customary fashion results in infringement of the claims of the ’893 Patent.

72. For example, each defendant explains to customers the benefits of using the Accused Instrumentalities, such as by touting their advantages of replicating data among multiple devices. Each defendant also induces its customers to use the Accused Instrumentalities to infringe other claims of the ’893 Patent. Each defendant specifically intended and was aware that the normal and customary use of the Accused Instrumentalities on compatible systems would infringe the ’893 Patent. Each defendant performed the acts that constitute induced infringement, and would induce actual infringement, with the knowledge of the ’893 Patent and with the knowledge, or willful blindness to the probability, that the induced acts would constitute infringement. On information and belief, each defendant engaged in such inducement to promote the sales of the Accused Instrumentalities, *e.g.*, through its user manuals, product support, marketing materials, demonstrations, installation support, and training materials to actively induce the users of the accused products to infringe the ’893 Patent. Accordingly, each defendant has induced and continues to induce end users of the accused products to use the accused products in their ordinary and customary way with compatible systems to make and/or use systems infringing the ’893 Patent, knowing that such use of the Accused Instrumentalities with compatible systems will result in infringement of the ’893 Patent. Accordingly, each defendant has been (since at least as of filing of the original complaint), and currently is, inducing infringement of the ’893 Patent, in violation of 35 U.S.C. § 271(b).

73. Each defendant has also infringed, and continues to infringe, claims of the ’893 Patent by offering to commercially distribute, commercially distributing, making, and/or importing the Accused Instrumentalities, which are used in practicing the process, or using the systems, of the ’893 Patent, and constitute a material part of the invention. Defendant knows the components in the Accused Instrumentalities to be especially made or especially adapted for use in infringement of the ’893 Patent, not a staple article, and

not a commodity of commerce suitable for substantial noninfringing use. For example, the ordinary way of using the Accused Instrumentalities infringes the patent claims, and as such, is especially adapted for use in infringement. Accordingly, each defendant has been, and currently is, contributorily infringing the '893 Patent, in violation of 35 U.S.C. § 271(c).

74. For similar reasons, each defendant also infringes the '893 Patent by supplying or causing to be supplied in or from the United States all or a substantial portion of the components of the Accused Instrumentalities, where such components are uncombined in whole or in part, in such manner as to actively induce the combination of such components outside of the United States in a manner that would infringe the '893 Patent if such combination occurred within the United States. For example, each defendant supplies or causes to be supplied in or from the United States all or a substantial portion of the hardware (e.g., Data Domain servers) and software (e.g., Data Domain OS) components of the Accused Instrumentalities in such a manner as to actively induce the combination of such components (e.g., by instructing users to combine multiple Data Domain servers into an infringing system) outside of the United States.

75. Each defendant also indirectly infringes the '893 Patent by supplying or causing to be supplied in or from the United States components of the Accused Instrumentalities that are especially made or especially adapted for use in infringing the '893 Patent and are not a staple article or commodity of commerce suitable for substantial non-infringing use, and where such components are uncombined in whole or in part, knowing that such components are so made or adapted and intending that such components are combined outside of the United States in a manner that would infringe the '893 Patent if such combination occurred within the United States. Because the Accused Instrumentalities are designed to operate as the claimed system and apparatus, the Accused Instrumentalities have no substantial non-infringing uses, and any other uses would be unusual, far-fetched, illusory, impractical, occasional, aberrant, or experimental. For example, each defendant supplies or causes to be supplied in or from the United States all

or a substantial portion of the hardware (e.g., separate Data Domain servers) and software (e.g., Data Domain OS) components that are especially made or especially adapted for use in the Accused Instrumentalities, where such hardware and software components are not staple articles or commodities of commerce suitable for substantial noninfringing use, knowing that such components are so made or adapted and intending that such components are combined outside of the United States, as evidenced by each defendant's own actions or instructions to users in, e.g., combining multiple Data Domain servers into infringing systems, and enabling and configuring the infringing functionalities of the Accused Instrumentalities.

76. As a result of Defendant's infringement of the '893 Patent, Plaintiff Data Scape is entitled to monetary damages in an amount adequate to compensate for each Defendant's infringement, but in no event less than a reasonable royalty for the use made of the invention by each Defendant, together with interest and costs as fixed by the Court.

**PRAYER FOR RELIEF**

WHEREFORE, Plaintiff Data Scape respectfully requests that this Court enter:

a. A judgment in favor of Plaintiff that each Defendant has infringed, either literally and/or under the doctrine of equivalents, the '581 Patent, '929 Patent, the '537 Patent, and the '893 Patent (collectively, "asserted patents");

b. A permanent injunction prohibiting each Defendant from further acts of infringement of the asserted patents;

c. A judgment and order requiring each Defendant to pay Plaintiff its damages, costs, expenses, and prejudgment and post-judgment interest for its infringement of the asserted patents, as provided under 35 U.S.C. § 284;

d. A judgment and order requiring each Defendant to provide an accounting and to pay supplemental damages to Data Scape, including without limitation, prejudgment and post-judgment interest;

e. A judgment and order finding that this is an exceptional case within the meaning of 35 U.S.C. § 285 and awarding to Plaintiff its reasonable attorneys' fees against each Defendant; and

f. Any and all other relief as the Court may deem appropriate and just under the circumstances.

**DEMAND FOR JURY TRIAL**

Plaintiff, under Rule 38 of the Federal Rules of Civil Procedure, requests a trial by jury of any issues so triable by right.

Dated: December 27, 2018

Respectfully submitted,

*/s/ Reza Mirzaie*

RUSS AUGUST & KABAT
Marc A. Fenster, SBN 181067
Email: mfenster@raklaw.com
Reza Mirzaie (CA SBN 246953)
Email: rmirzaie@raklaw.com
Brian D. Ledahl (CA SBN 186579)
Email: bledahl@raklaw.com
Paul A. Kroeger (CA SBN 229074)
Email: pkroeger@raklaw.com
C. Jay Chung (CA SBN 252794)
Email: jchung@raklaw.com
Philip X. Wang (CA SBN 262239)
Email: pwang@raklaw.com
12424 Wilshire Boulevard, 12th Floor
Los Angeles, California 90025
Telephone: (310) 826-7474
Facsimile: (310) 826-6991

*Attorneys for Plaintiff Data Scape Limited*